arXiv:1604.06637v3 [stat.ME] 29 Aug 2016

Robust and Sparse Regression via γ-divergence

Takayuki Kawashima[1] and Hironori Fujisawa[1,2,3]

[1] Department of Statistical Science, The Graduate University for Advanced Studies, Tokyo
E-mail: t-kawa@ism.ac.jp
[2]The Institute of Statistical Mathematics, Tokyo
E-mail: fujisawa@ism.ac.jp
[3] Department of Mathematical Statistics, Nagoya University Graduate School of Medicine

Abstract

In high-dimensional data, many sparse regression methods have been proposed. However, they may not be robust against outliers. Recently, the use of density power weight has been studied for robust parameter estimation and the corresponding divergences have been discussed. One of such divergences is the γ-divergence and the robust estimator using the γ-divergence is known for having a strong robustness. In this paper, we consider the robust and sparse regression based on γ-divergence. We extend the γ-divergence to the regression problem and show that it has a strong robustness under heavy contamination even when outliers are heterogeneous. The loss function is constructed by an empirical estimate of the γ-divergence with sparse regularization and the parameter estimate is defined as the minimizer of the loss function. To obtain the robust and sparse estimate, we propose an efficient update algorithm which has a monotone decreasing property of the loss function. Particularly, we discuss a linear regression problem with L_1 regularization in detail. In numerical experiments and real data analyses, we see that the proposed method outperforms past robust and sparse methods.

1 Introduction

In high-dimensional data, sparse regression methods have been intensively studied. The Lasso (Tibshirani 1996) is a typical sparse linear regression method with L_1 regularization, but is not robust against outliers. Recently, robust and sparse linear regression methods have been proposed. The robust least angle regression (RLARS)(Khan et al. 2007) is a robust version of LARS (Efron et al. 2004), which replaces the sample correlation by a robust estimate of correlation in the update algorithm. The sparse least trimmed squares (sLTS)(Alfons et al. 2013) is a sparse version of well-known robust linear regression method LTS (Rousseeuw 1984) based on the trimmed loss function with L_1 regularization.

Recently, the robust parameter estimation using density power weight has been discussed by Basu et al. (1998), Basu et al. (2010), Jones et al. (2001), Windham (1995), Fujisawa and Eguchi (2008), Kanamori and Fujisawa (2015), and so on. The density power weight gives a small weight to the terms related to outliers and then the parameter estimation becomes robust against outliers. Among them, the γ-divergence proposed by Fujisawa and Eguchi (2008) is known for having a strong robustness, which implies that the latent bias can be sufficiently small even under heavy contamination. In addition, to obtain the robust estimate, an efficient update algorithm was proposed with a monotone decreasing property of the loss function.

In this paper, we propose the robust and sparse regression problem based on γ-divergence. First, we extend the γ-divergence to regression problem and show a strong robustness under heavy contamination even when outliers are heterogeneous. Next, we consider a loss function based on the γ-divergence with sparse regularization and propose an update algorithm to obtain the robust and sparse estimate. Fujisawa and Eguchi (2008) used a Pythagorian relation on the γ-divergence, but it is not compatible with sparse regularization. Instead of this relation, we use the Majorization-Minimization algorithm (Hunter and Lange 2004). This idea is deeply considered in a linear regression problem with L_1 regularization. Finally, in numerical experiments and real data analyses, we show that our method outperforms other robust and sparse methods.

2 Regression based on γ-divergence

The γ-divergence was defined for two probability density functions and its properties were investigated by Fujisawa and Eguchi (2008). In this section, the γ-divergence is extended to the regression problem, in other words, de-

fined for two conditional probability density functions and the corresponding robust properties are presented.

2.1 γ-divergence for regression

We suppose $g(x,y)$, $g(y|x)$ and $g(x)$ are the underlying probability density functions of (x,y), y given x and x, respectively. Let $f(y|x)$ be another parametric conditional probability density function of y given x. Let us define the γ-cross entropy for regression by

$$\begin{aligned}
&d_\gamma(g(y|x), f(y|x); g(x)) \\
&= -\frac{1}{\gamma}\log\int\left(\int g(y|x)f(y|x)^\gamma dy\right)g(x)dx + \frac{1}{1+\gamma}\log\int\left(\int f(y|x)^{1+\gamma}dy\right)g(x)dx \\
&= -\frac{1}{\gamma}\log\int\int f(y|x)^\gamma g(x,y)dxdy + \frac{1}{1+\gamma}\log\int\left(\int f(y|x)^{1+\gamma}dy\right)g(x)dx.
\end{aligned} \tag{2.1}$$

The γ-divergence for regression is defined by

$$D_\gamma(g(y|x), f(y|x); g(x)) = -d_\gamma(g(y|x), g(y|x); g(x)) + d_\gamma(g(y|x), f(y|x); g(x)). \tag{2.2}$$

The γ-divergence for regression was first proposed by Fujisawa and Eguchi (2008) and many properties were already shown. However, we adopt the definition (2.2), which is slightly different from the past one, because (2.2) satisfies the Pythagorian relation approximately (see Sec 2.3).

Theorem 2.1. *We can show that*

$$\begin{aligned}
&\text{(i)} \quad D_\gamma(g(y|x), f(y|x); g(x)) \geq 0, \\
&\text{(ii)} \quad D_\gamma(g(y|x), f(y|x); g(x)) = 0 \quad \Leftrightarrow \quad g(y|x) = f(y|x) \quad \text{(a.e.)}, \\
&\text{(iii)} \quad \lim_{\gamma\to 0} D_\gamma(g(y|x), f(y|x); g(x)) = \int D_{KL}(g(y|x), f(y|x))g(x)dx,
\end{aligned}$$

where $D_{KL}(g(y|x), f(y|x)) = \int g(y|x)\log g(y|x)dy - \int g(y|x)\log f(y|x)dy$.

The proof is in the Appendix. In what follows, we refer to the regression based on γ-divergence as the γ-regression.

2.2 Estimation for γ-regression

Let $f(y|x;\theta)$ be the conditional probability density function of y given x with parameter θ. The target parameter can be considered by

$$\begin{aligned}\theta^*_\gamma &= \operatorname*{argmin}_\theta D_\gamma(g(y|x), f(y|x;\theta); g(x)) \\ &= \operatorname*{argmin}_\theta d_\gamma(g(y|x), f(y|x;\theta); g(x)). \end{aligned} \tag{2.3}$$

When $g(y|x) = f(y|x;\theta^*)$, we have $\theta^*_\gamma = \theta^*$.

Let $(x_1, y_1), \ldots, (x_n, y_n)$ be the observations randomly drawn from the underlying distribution $g(x, y)$. Using the formula (2.1), the γ-cross entropy for regression, $d_\gamma(g(y|x), f(y|x;\theta); g(x))$, can be empirically estimated by

$$\begin{aligned}&\bar{d}_\gamma(f(y|x;\theta)) \\ &= -\frac{1}{\gamma}\log\left\{\frac{1}{n}\sum_{i=1}^n f(y_i|x_i;\theta)^\gamma\right\} + \frac{1}{1+\gamma}\log\left\{\frac{1}{n}\sum_{i=1}^n \int f(y|x_i;\theta)^{1+\gamma}dy\right\}.\end{aligned}$$

By virtue of (2.3), we define the γ-estimator by

$$\hat{\theta}_\gamma = \operatorname*{argmin}_\theta \bar{d}_\gamma(f(y|x;\theta)).$$

In a similar way to in Fujisawa and Eguchi (2008), we can show the consistency of $\hat{\theta}_\gamma$ to θ^*_γ under some conditions.

2.3 Robust properties for γ-regression

In this subsection, the robust properties are presented from two viewpoints of latent bias and Pythagorian relation. The latent bias was discussed in Fujisawa and Eguchi (2008) and Kanamori and Fujisawa (2015), which is described later. Using the results obtained there, the Pythagorian relation is shown in Theorems 2.2 and 2.3.

Let $f^*(y|x) = f_{\theta^*}(y|x) = f(y|x;\theta^*)$ and $\delta(y|x)$ be the target conditional probability density function and the contamination conditional probability density function related to outliers, respectively. Let ϵ and $\epsilon(x)$ denote the outlier ratios which are independent and dependent of x, respectively. Under homogeneous and heterogeneous contaminations, we suppose that the underlying conditional probability density function can be expressed as

$$\begin{aligned}g(y|x) &= (1-\epsilon)f(y|x;\theta^*) + \epsilon\delta(y|x), \\ g(y|x) &= (1-\epsilon(x))f(y|x;\theta^*) + \epsilon(x)\delta(y|x).\end{aligned}$$

Let

$$\nu_{f,\gamma}(x) = \left\{ \int \delta(y|x) f(y|x)^\gamma dy \right\}^{\frac{1}{\gamma}} \qquad (\gamma > 0).$$

and let

$$\nu_{f,\gamma} = \left\{ \int \nu_{f,\gamma}(x)^\gamma g(x) dx \right\}^{\frac{1}{\gamma}}.$$

Here we assume that

$$\nu_{f_{\theta^*},\gamma} \approx 0,$$

which implies that $\nu_{f_{\theta^*},\gamma}(x) \approx 0$ for any x (a.e.) and illustrates that the contamination conditional probability density function $\delta(y|x)$ lies on the tail of the target conditional probability density function $f(y|x;\theta^*)$. For example, if $\delta(y|x)$ is the dirac function at the outlier $y_\dagger(x)$ given x, then we have $\nu_{f_{\theta^*},\gamma}(x) = f(y_\dagger(x)|x;\theta^*)$, which should be sufficiently small because $y_\dagger(x)$ is an outlier. In this subsection, we show that $\theta^*_\gamma - \theta^*$ is expected to be small even if ϵ or $\epsilon(x)$ is not small. To make the discussion easier, we prepare the monotone transformation of the γ-cross entropy for regression by

$$\begin{aligned}
\tilde{d}_\gamma(g(y|x), f(y|x;\theta); g(x)) &= -\exp\left\{-\gamma d_\gamma(g(y|x), f(y|x;\theta); g(x))\right\} \\
&= -\frac{\int \left(\int g(y|x) f(y|x;\theta)^\gamma dy \right) g(x) dx}{\left\{ \int \left(\int f(y|x;\theta)^{1+\gamma} dy \right) g(x) dx \right\}^{\frac{\gamma}{1+\gamma}}}.
\end{aligned}$$

2.3.1 Homogeneous contamination

Here we provide the following proposition, which was given in Kanamori and Fujisawa (2015).

Proposition 2.1.

$$\begin{aligned}
&\tilde{d}_\gamma(g(y|x), f(y|x;\theta); g(x)) \\
&= (1-\epsilon)\tilde{d}_\gamma(f(y|x;\theta^*), f(y|x;\theta); g(x)) - \frac{\epsilon \nu_{f_\theta,\gamma}^\gamma}{\left\{ \int \left(\int f(y|x;\theta)^{1+\gamma} dy \right) g(x) dx \right\}^{\frac{\gamma}{1+\gamma}}}.
\end{aligned}$$

Recall that θ^*_γ and θ^* are also the minimizers of $\tilde{d}_\gamma(g(y|x), f(y|x;\theta); g(x))$ and $\tilde{d}_\gamma(f(y|x;\theta^*), f(y|x;\theta); g(x))$, respectively. We can expect $\nu_{f_\theta,\gamma} \approx 0$ from

the assumption $\nu_{f_{\theta^*},\gamma} \approx 0$ if the tail behavior of $f(y|x;\theta)$ is close to that of $f(y|x;\theta^*)$. We see from Proposition 2.1 and the condition $\nu_{f_\theta,\gamma} \approx 0$ that

$$\begin{aligned}
\theta^*_\gamma &= \operatorname*{argmin}_\theta \tilde{d}_\gamma(g(y|x), f(y|x;\theta); g(x)) \\
&= \operatorname*{argmin}_\theta \left[(1-\epsilon)\tilde{d}_\gamma(f(y|x;\theta^*), f(y|x;\theta); g(x)) - \frac{\epsilon \nu^\gamma_{f_\theta,\gamma}}{\left\{ \int \left(\int f(y|x;\theta)^{1+\gamma} dy \right) g(x) dx \right\}^{\frac{\gamma}{1+\gamma}}} \right] \\
&\approx \operatorname*{argmin}_\theta (1-\epsilon)\tilde{d}_\gamma(f(y|x;\theta^*), f(y|x;\theta); g(x)) \\
&= \theta^*.
\end{aligned}$$

Therefore, under homogeneous contamination, it can be expected that the latent bias $\theta^*_\gamma - \theta^*$ is small even if ϵ is not small. Moreover, we can show the following theorem, using Proposition 2.1.

Theorem 2.2. *Let $\nu = \max\{\nu_{f_\theta,\gamma}, \nu_{f_{\theta^*},\gamma}\}$. Then, the Pythagorean relation among $g(y|x)$, $f(y|x;\theta^*)$, $f(y|x;\theta)$ approximately holds:*

$$\begin{aligned}
&D_\gamma(g(y|x), f(y|x;\theta); g(x)) - D_\gamma(g(y|x), f(y|x;\theta^*); g(x)) - D_\gamma(f(y|x;\theta^*), f(y|x;\theta); g(x)) \\
&= O(\nu^\gamma).
\end{aligned}$$

The proof is in the Appendix. The Pythagorean relation implies that the minimization of the divergence from $f(y|x;\theta)$ to the underlying conditional probability density function $g(y|x)$ is approximately the same as that to the target conditional probability density function $f(y|x;\theta^*)$. Therefore, under homogeneous contamination, we can see why our proposed method works well in terms of minimization of γ-divergence.

2.3.2 Heterogeneous contamination

Under heterogeneous contamination, we assume that the parametric conditional probability density function $f(y|x;\theta)$ is a location-scale family given by

$$f(y|x;\theta) = \frac{1}{\sigma} s\left(\frac{y - q(x;\xi)}{\sigma} \right),$$

where $s(y)$ is a probability density function, σ is a scale parameter and $q(x;\xi)$ is a location function with a regression parameter ξ, e.g., $q(x;\xi) = \xi^T x$. Then, we can obtain

$$\begin{aligned}
\int f(y|x;\theta)^{1+\gamma} dy &= \int \frac{1}{\sigma^{1+\gamma}} s\left(\frac{y - q(x;\xi)}{\sigma} \right)^{1+\gamma} dy \\
&= \sigma^{-\gamma} \int s(z)^{1+\gamma} dz,
\end{aligned}$$

That does not depend on the explanatory variable x. Here, we provide the following proposition, which was given in Kanamori and Fujisawa (2015).

Proposition 2.2.

$$\begin{aligned}
&\tilde{d}_\gamma(g(y|x), f(y|x;\theta); g(x)) \\
&= \left(1 - \int \epsilon(x)g(x)dx\right)^{\frac{\gamma}{1+\gamma}} \tilde{d}_\gamma(f(y|x;\theta^*), f(y|x;\theta); (1-\epsilon(x))g(x)) \\
&\qquad - \frac{\int \nu_{f_\theta,\gamma}(x)^\gamma \epsilon(x) g(x) dx}{\left\{\sigma^{-\gamma}\int s(z)^{1+\gamma}dz\right\}^{\frac{\gamma}{1+\gamma}}}.
\end{aligned}$$

The second term $\frac{\int \nu_{f_\theta,\gamma}(x)^\gamma \epsilon(x) g(x) dx}{\left\{\sigma^{-\gamma}\int s(z)^{1+\gamma}dz\right\}^{\frac{\gamma}{1+\gamma}}}$ can be approximated to be 0 from the condition $\nu_{f_\theta,\gamma} \approx 0$ and $\epsilon(x) < 1$ as follows:

$$\begin{aligned}
\frac{\int \nu_{f_\theta,\gamma}(x)^\gamma \epsilon(x) g(x) dx}{\left\{\sigma^{-\gamma}\int s(z)^{1+\gamma}dz\right\}^{\frac{\gamma}{1+\gamma}}} &< \frac{\int \nu_{f_\theta,\gamma}(x)^\gamma g(x) dx}{\left\{\sigma^{-\gamma}\int s(z)^{1+\gamma}dz\right\}^{\frac{\gamma}{1+\gamma}}} \\
&= \frac{\nu_{f_\theta,\gamma}^\gamma}{\left\{\sigma^{-\gamma}\int s(z)^{1+\gamma}dz\right\}^{\frac{\gamma}{1+\gamma}}} \\
&\approx 0 \qquad (2.4)
\end{aligned}$$

We see from Proposition 2.2 and (2.4) that

$$\theta^*_\gamma = \operatorname*{argmin}_\theta \tilde{d}_\gamma(g(y|x), f(y|x;\theta); g(x))$$

$$\begin{aligned}
&= \operatorname*{argmin}_\theta \left[\left(1 - \int \epsilon(x)g(x)dx\right)^{\frac{\gamma}{1+\gamma}} \tilde{d}_\gamma(f(y|x;\theta^*), f(y|x;\theta); (1-\epsilon(x))g(x))\right. \\
&\qquad \left. - \frac{\int \nu_{f_\theta,\gamma}(x)^\gamma \epsilon(x) g(x) dx}{\left\{\sigma^{-\gamma}\int s(z)^{1+\gamma}dz\right\}^{\frac{\gamma}{1+\gamma}}}\right] \\
&\approx \operatorname*{argmin}_\theta \left(1 - \int \epsilon(x)g(x)dx\right)^{\frac{\gamma}{1+\gamma}} \tilde{d}_\gamma(f(y|x;\theta^*), f(y|x;\theta); (1-\epsilon(x))g(x)) \\
&= \theta^*.
\end{aligned}$$

Therefore, under heterogeneous contamination in a location-scale family, it can be expected that the latent bias $\theta^*_\gamma - \theta^*$ is small even if $\epsilon(x)$ is not small. Moreover, we can show the following theorem, using Proposition 2.2.

Theorem 2.3. *Let* $\nu = \max\{\nu_{f_\theta,\gamma}, \nu_{f_{\theta^*},\gamma}\}$. *Then, the following relation among* $g(y|x)$, $f(y|x;\theta^*)$, $f(y|x;\theta)$ *approximately holds:*

$$
\begin{aligned}
&D_\gamma(g(y|x), f(y|x;\theta); g(x)) \\
&- D_\gamma(g(y|x), f(y|x;\theta^*); g(x)) - D_\gamma(f(y|x;\theta^*), f(y|x;\theta); (1-\epsilon(x))g(x)) \\
&= O(\nu^\gamma).
\end{aligned}
$$

The proof is in the Appendix. The above is slightly different from a conventional Pythagorean relation, because the base measure changes from $g(x)$ to $(1-\epsilon(x))g(x)$ in part. But, it also implies that the minimization of the divergence from $f(y|x;\theta)$ to the underlying conditional probability density function $g(y|x)$ is approximately the same as that to the target conditional probability density function $f(y|x;\theta^*)$. Therefore, under heterogeneous contamination in a location-scale family, we can see why our proposed method works well in terms of minimization of γ-divergence.

3 Sparse γ-regression

3.1 Estimation for sparse γ-regression

The empirical estimation of the γ-cross entropy with a penalty term can be given by

$$L_\gamma(\theta;\lambda) = \bar{d}_\gamma(f(y|x;\theta)) + \lambda P(\theta),$$

where $P(\theta)$ is a penalty for parameter θ and λ is a tuning parameter for the penalty term. As an example of penalty term, we consider L_1 (Lasso, Tibshirani 1996), elasticnet (Zou and Hastie 2005), and so on. The sparse γ-estimator can be proposed by

$$\hat{\theta}_S = \operatorname*{argmin}_\theta L_\gamma(\theta;\lambda). \tag{3.1}$$

To obtain the minimizer, we propose the iterative algorithm by Majorization-Minimization algorithm (MM algorithm) (Hunter and Lange 2004).

3.2 MM algorithm for sparse γ-regression

The MM algorithm is constructed as follows. Let $h(\eta)$ be the objective function. Let us prepare the majorization function h_{MM} satisfying

$$
\begin{aligned}
h_{MM}(\eta^{(m)}|\eta^{(m)}) &= h(\eta^{(m)}), \\
h_{MM}(\eta|\eta^{(m)}) &\geq h(\eta) \quad \text{for all } \eta,
\end{aligned}
$$

where $\eta^{(m)}$ is the parameter of the m-th iterative step for $m = 0, 1, 2, \ldots$. Let us consider the iterative algorithm by

$$\eta^{(m+1)} = \operatorname*{argmin}_{\eta} h_{MM}(\eta|\eta^{(m)}).$$

Then, we can show that the objective function $h(\eta)$ monotonically decreases at each step, because

$$h(\eta^{(m)}) = h_{MM}(\eta^{(m)}|\eta^{(m)}) \geq h_{MM}(\eta^{(m+1)}|\eta^{(m)}) \geq h(\eta^{(m+1)}).$$

Note that $\eta^{(m+1)}$ is not necessary to be minimizer of $h_{MM}(\eta|\eta^{(m)})$. We only need $h_{MM}(\eta^{(m)}|\eta^{(m)}) \geq h_{MM}(\eta^{(m+1)}|\eta^{(m)})$.

We construct the majorization function for the sparse γ-regression by the following inequality:

$$\kappa(z^T\eta) \leq \sum_i \frac{z_i\eta_i^{(m)}}{z^T\eta^{(m)}}\kappa\left[\eta_i\frac{z^T\eta^{(m)}}{\eta_i^{(m)}}\right], \tag{3.2}$$

where $\kappa(u)$ is a convex function, $z = (z_1, \ldots, z_n)^T$, $\eta = (\eta_1, \ldots, \eta_n)^T$, $\eta^{(m)} = (\eta_1^{(m)}, \ldots, \eta_n^{(m)})^T$ and z_i, η_i and $\eta_i^{(m)}$ are positive. The inequality (3.2) holds from Jensen's inequality. Here, we take $z_i = \frac{1}{n}$, $\eta_i = f(y_i|x_i;\theta)^\gamma$, $\eta_i^{(m)} = f(y_i|x_i;\theta^{(m)})^\gamma$ and $\kappa(u) = -\log u$ in (3.2). We can propose the majorization function as follows:

$$\begin{aligned}
&h(\theta)\\
&= L_\gamma(\theta;\lambda)\\
&= -\frac{1}{\gamma}\log\left\{\frac{1}{n}\sum_{i=1}^n f(y_i|x_i;\theta)^\gamma\right\} + \frac{1}{1+\gamma}\log\left\{\frac{1}{n}\sum_{i=1}^n\int f(y|x_i;\theta)^{1+\gamma}dy\right\} + \lambda P(\theta)\\
&\leq -\frac{1}{\gamma}\sum_{i=1}^n \frac{\frac{1}{n}f(y_i|x_i;\theta^{(m)})^\gamma}{\frac{1}{n}\sum_{l=1}^n f(y_l|x_l;\theta^{(m)})^\gamma}\log\left\{f(y_i|x_i;\theta)^\gamma\frac{\frac{1}{n}\sum_{l=1}^n f(y_l|x_l;\theta^{(m)})^\gamma}{f(y_i|x_i;\theta^{(m)})^\gamma}\right\}\\
&\qquad + \frac{1}{1+\gamma}\log\left\{\frac{1}{n}\sum_{i=1}^n\int f(y|x_i;\theta)^{1+\gamma}dy\right\} + \lambda P(\theta)\\
&= -\sum_{i=1}^n \alpha_i^{(m)}\log f(y_i|x_i;\theta) + \frac{1}{1+\gamma}\log\left\{\frac{1}{n}\sum_{i=1}^n\int f(y|x_i;\theta)^{1+\gamma}dy\right\}\\
&\qquad + \lambda P(\theta) + const\\
&= h_{MM}(\theta|\theta^{(m)}) + const,
\end{aligned}$$

where $\alpha_i^{(m)} = f(y_i|x_i;\theta^{(m)})^\gamma \Big/ \sum_{l=1}^n f(y_l|x_l;\theta^{(m)})^\gamma$ and *const* is a term which does not depend on the parameter θ.

The first term on the original target function $h(\theta)$ is a mixture type of densities, which is not easy to be optimized, while the first term on $h_{MM}(\theta|\theta^{(m)})$ is a weighted log-likelihood, which is often easy to be optimized.

3.3 Sparse γ-linear regression

Let $f(y|x;\theta)$ be the conditional density with $\theta = (\beta_0, \beta, \sigma^2)$, given by

$$f(y|x;\theta) = \phi(y; \beta_0 + x^T\beta, \sigma^2),$$

where $\phi(y;\mu,\sigma^2)$ is the normal density with mean parameter μ and variance parameter σ^2. Suppose that $P(\theta)$ is the L_1 regularization $||\beta||_1$. Then, after a simple calculation, we have

$$h_{MM}(\theta|\theta^{(m)}) = \frac{1}{2(1+\gamma)}\log\sigma^2 + \frac{1}{2}\sum_{i=1}^n \alpha_i^{(m)}\frac{(y_i - \beta_0 - x_i^T\beta)^2}{\sigma^2} + \lambda||\beta||_1. \tag{3.3}$$

This function is easy to be optimized by an update algorithm. For a fixed value of σ^2, the function h_{MM} is almost the same as Lasso except for the weight, so that it can be updated using the coordinate decent algorithm with a decreasing property of the loss function. For a fixed value of $(\beta_0, \beta^T)^T$, the function h_{MM} is easy to be minimized. Consequently, the update formula can be obtained by

$$\beta_0^{(m+1)} = \sum_{i=1}^n \alpha_i^{(m)}(y_i - {x_i}^T\beta^{(m)}), \tag{3.4}$$

$$\beta_j^{(m+1)} = S\left(\sum_{i=1}^n \alpha_i^{(m)}(y_i - \beta_0^{(m+1)} - r_i^{[-j]})x_{ij}, {\sigma^2}^{(m)}\lambda\right) \Big/ \left(\sum_{i=1}^n \alpha_i^{(m)}x_{ij}^2\right) \quad (j = 1,\ldots,p), \tag{3.5}$$

$${\sigma^2}^{(m+1)} = (1+\gamma)\sum_{i=1}^n \alpha_i^{(m)}(y_i - \beta_0^{(m+1)} - x_i^T\beta^{(m+1)})^2, \tag{3.6}$$

where $S(t,\lambda) = \mathrm{sign}(t)(|t| - \lambda)_+$ and $r_i^{[-j]} = \sum_{k\neq j} x_{ik}\beta_k$, and the decreasing property holds:

$$h_{MM}(\theta^{(m+1)}|\theta^{(m)}) \leq h_{MM}(\theta^{(m)}|\theta^{(m)}). \tag{3.7}$$

It should be noted that h_{MM} is convex with respect to parameter β_0, β and has the global minimum with respect to parameter σ^2, but not convex with respect to some of them, so that the starting value of update formula is important. This issue is often important in robust statistics and discussed in Sec 4.4.

Moreover, we explain how to implement coordinate decent algorithm. In practice, we also use active set strategy (Friedman et al. 2007). Specifically, for a given $\beta^{(m)}$, we only update the non-zero coordinates of $\beta^{(m)}$ in the active set, until they are converged. Then, the non-active set parameter estimates are updated once. When they remain zero, the coordinate descent algorithm stops.

3.4 Robust cross-validation

In sparse regression, a regularization parameter is often selected via some criterion. Cross-validation is often used for selecting the regularization parameter. However, the ordinal cross-validation will fail due to outliers. Therefore, we propose the robust cross-validation based on the γ-cross entropy. Let $\hat{\theta}_\gamma$ be the robust estimate based on the γ-cross entropy. The ordinal cross validation is based on the squared error and it can also be constructed using the KL-cross entropy with normal density. The cross-validation based on the γ-cross entropy can be given by

$$\begin{aligned}\mathrm{RoCV}(\lambda) = &-\frac{1}{\gamma_0}\log\left\{\frac{1}{n}\sum_{i=1}^{n} f(y_i|x_i;\hat{\theta}_\gamma^{[-i]})^{\gamma_0}\right\} \\ &+\frac{1}{1+\gamma_0}\log\left\{\frac{1}{n}\sum_{i=1}^{n}\int f(y|x_i;\hat{\theta}_\gamma^{[-i]})^{1+\gamma_0}dy\right\},\end{aligned}$$

where $\hat{\theta}_\gamma^{[-i]}$ is the γ-estimator deleting the i-th observation and γ_0 is an appropriate tuning parameter. We can also adopt the K-fold cross validation to reduce the computational task (Hastie et al. 2010).

Here, we give a small modification on the above. We often focus only on the mean structure for prediction, not on the variance parameter. Therefore, in this paper, $\hat{\theta}_\gamma^{[-i]}$ is replaced by $(\hat{\beta}_\gamma^{[-i]}, \hat{\sigma}_\gamma^2)$.

3.5 Redescending Property

First, we review a redescending property on M-estimation (see, e.g., Maronna et al. 2006), which is often used in robust statistics. Suppose that the estimating equation is given by $\sum_{i=1}^{n}\zeta(z_i;\theta)=0$. Let $\hat{\theta}$ be a solution of the estimating

equation. The bias caused by outlier z_o is expressed as $\hat{\theta}_{n=\infty} - \theta^*$, where $\hat{\theta}_{n=\infty}$ is the limiting value of $\hat{\theta}$ and θ^* is the true parameter. We hope the bias is small even if the outlier z_o exists. Under some conditions, the bias can be approximated to $\epsilon \mathrm{IF}(z_o; \theta^*)$, where ϵ is a small outlier ratio and $\mathrm{IF}(z; \theta^*)$ is the influence function. The bias is expected to be small when the influence function is small. The influence function can be expressed as $\mathrm{IF}(z; \theta^*) = A\zeta(z; \theta^*)$, where A is a matrix independent of z, so that the bias is also expected to be small when $\zeta(z_o; \theta^*)$ is small. In particular, the estimating equation is said to have a redescending property if $\zeta(z; \theta^*)$ goes to zero as $||z||$ goes to infinity. This property is favorable in robust statistics, because the bias is expected to be sufficiently small when z_o is very large.

Here, we prove a redescending property on the sparse γ-linear regression, i.e., when $f(y|x;\theta) = \phi(y; \beta_0 + x^T\beta, \sigma^2)$ with $\theta = (\beta_0, \beta, \sigma^2)$ for fixed x. Recall that the estimate of the sparse γ-linear regression is the minimizer of the loss function

$$L_\gamma(\theta; \lambda) = -\frac{1}{\gamma} \log \left\{ \frac{1}{n} \sum_{i=1}^{n} \phi(y_i; \beta_0 + {x_i}^T\beta, \sigma^2)^\gamma \right\} + b_\gamma(\theta; \lambda),$$

where

$$b_\gamma(\theta; \lambda) = \frac{1}{1+\gamma} \log \left\{ \frac{1}{n} \sum_{i=1}^{n} \int \phi(y; \beta_0 + {x_i}^T\beta, \sigma^2)^{1+\gamma} dy \right\} + \lambda ||\beta||_1.$$

Then, the estimating equation is given by

$$\begin{aligned} 0 &= \frac{\partial}{\partial \theta} L_\gamma(\theta; \lambda) \\ &= -\frac{\sum_{i=1}^{n} \phi(y_i; \beta_0 + {x_i}^T\beta, \sigma^2)^\gamma s(y_i|x_i;\theta)}{\sum_{i=1}^{n} \phi(y_i; \beta_0 + {x_i}^T\beta, \sigma^2)^\gamma} + \frac{\partial}{\partial \theta} b_\gamma(\theta; \lambda), \end{aligned}$$

where

$$s(y|x;\theta) = \frac{\partial \log \phi(y; \beta_0 + x^T\beta, \sigma^2)}{\partial \theta}.$$

This can be expressed by the M-estimation formula given by

$$0 = \sum_{i=1}^{n} \psi(y_i|x_i;\theta), \tag{3.8}$$

where

$$\psi(y|x;\theta) = \phi(y; \beta_0 + x^T\beta, \sigma^2)^\gamma s(y|x;\theta) - \phi(y; \beta_0 + x^T\beta, \sigma^2)^\gamma \frac{\partial}{\partial \theta} b_\gamma(\theta; \lambda).$$

We can easily show that as $||y||$ goes to infinity, $\phi(y; \beta_0 + x^T\beta, \sigma^2)$ goes to zero and $\phi(y; \beta_0 + x^T\beta, \sigma^2)s(y|x; \theta)$ also goes to zero. Therefore, the function $\psi(y|x; \theta)$ goes to zero as $||y||$ goes to infinity, so that the estimating equation has a redescending property.

4 Numerical experiment

In this section, we compare our method (Sparse γ-linear regression) with the representative sparse linear regression method, *Least absolute shrinkage and selection operator* (Lasso)(Tibshirani 1996), and the robust and sparse regression methods, *sparse least trimmed squares* (sLTS)(Alfons et al. 2013) and *robust least angle regression* (RLARS)(Khan et al. 2007).

4.1 Regression models for simulation

We used the simulation model given by

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \cdots + \beta_p x_p + e, \quad e \sim N(0, 0.5^2).$$

The sample size and the number of explanatory variables were set to be $n = 100$ and $p = 100, 200$, respectively. The true coefficients were given by

$$\begin{aligned} &\beta_1 = 1,\ \beta_2 = 2,\ \beta_4 = 4,\ \beta_7 = 7,\ \beta_{11} = 11, \\ &\beta_j = 0 \text{ for } j \in \{0, \ldots, p\} \backslash \{1, 2, 4, 7, 11\}. \end{aligned} \tag{4.1}$$

We arranged a broad range of regression coefficients to observe sparsity for various degrees of regression coefficients. The explanatory variables were generated from a normal distribution $N(0, \Sigma)$, where the covariance matrix $\Sigma = (\sigma_{ij})_{1 \le i,j \le p}$ was given by $\sigma_{ij} = \rho^{|i-j|}$. We generated 100 random samples.

Outliers were incorporated into simulations. We investigated two cases for outlier ratio ($\epsilon = 0.1$ and 0.3) and two cases for outlier pattern: (a) The outliers were generated around the middle part of the explanatory variable, where the explanatory variables were generated from $N(0, 0.5^2)$ and the error terms were generated from $N(20, 0.5^2)$. (b) The outliers were generated around the edge part of the explanatory variable, where the explanatory variables were generated from $N(-1.5, 0.5^2)$ and the error terms were generated from $N(20, 0.5^2)$.

4.2 Performance measure

The root mean squared prediction error (RMSPE) and mean squared error (MSE) were examined to verify predictive performance and fitness of regres-

sion coefficient:

$$\mathrm{RMSPE}(\hat{\beta}) = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i^* - x_i^{*T}\hat{\beta})^2},$$

$$\mathrm{MSE} = \frac{1}{p+1}\sum_{j=0}^{p}(\beta_j^* - \hat{\beta}_j)^2,$$

where y_i^* and $x_i^*(i = 1, \ldots, n)$ is the test sample generated from the simulation model without outliers and β_j^*'s are the true coefficients. The true positive rate (TPR) and true negative rate (TNR) were also reported to verify the sparsity:

$$\mathrm{TPR}(\hat{\beta}) = \frac{|\{j \in \{1, \ldots, p\} : \hat{\beta}_j \neq 0 \wedge \beta_j^* \neq 0\}|}{|\{j \in \{1, \ldots, p\} : \beta_j^* \neq 0\}|},$$

$$\mathrm{TNR}(\hat{\beta}) = \frac{|\{j \in \{1, \ldots, p\} : \hat{\beta}_j = 0 \wedge \beta_j^* = 0\}|}{|\{j \in \{1, \ldots, p\} : \beta_j^* = 0\}|}.$$

4.3 Comparative methods

In this subsection, we explain three comparative methods; Lasso, RLARS, sLTS.

The Lasso is performed by R-package "glmnet". The regularization parameter λ_{lasso} is selected by grid search via cross-validation in "glmnet". We used "glmnet" by default.

The RLARS is performed by R-package "robustHD". This is a robust version of LARS (Efron et al. 2004). The optimal model is selected via BIC by default.

The sLTS is performed by R-package "robustHD". The sLTS has the regularization parameter λ_{sLTS} and the fraction parameter α of squared residuals used for trimmed squares. First, the regularization parameter λ_{sLTS} is selected by grid search via BIC. The number of grids is 40 by default. However, we considered that this would be small under heavy contamination. Therefore, we used 80 grids under heavy contamination. Next, the fraction parameter α is 0.75 by default. In the case of $\alpha = 0.75$, the ratio of outlier is less than 25%. We considered this would be small under heavy contamination and large under low contamination in terms of statistical efficiency. Therefore, we used 0.65, 0.75, 0.85 as α under low contamination and 0.50, 0.65, 0.75 under heavy contamination.

4.4 Details of our method

4.4.1 Initial points

In our method, we need an initial point to obtain the estimate, because we use the iterative algorithm proposed in Sec 3.3. The estimate of other conventional robust and sparse regression methods would give a good initial point. For another choice, the estimate of the RANSAC (random sample consensus) algorithm would also give a good initial point. In this experiment, we used the estimate of the sLTS as an initial point.

4.4.2 How to choose tuning parameters

In our method, we have to choose some tuning parameters. The parameter γ in γ-divergence was set to 0.1 or 0.5. The parameter γ_0 in robust cross-validation was set to 0.5. In our experience, the result via Rocv is not sensitive to the selection of γ_0 when γ_0 is enough large, e.g. $\gamma_0 = 0.5, 1$. The parameter λ of L_1 regularization is often selected via grid search. We used 50 grids on range $[0.05\lambda_0, \lambda_0]$ with log scale, where λ_0 is an estimate of λ which would shrink regression coefficients to zero.

4.5 Result

Table 1 is the low contamination case with outlier pattern (a). For the RMSPE, our method outperformed other comparative methods (the oracle value of the RMSPE is 0.5). For the TPR and TNR, the sLTS showed a similar performance to our method. The Lasso presented the worst performance, because it is sensitive to outliers. Table 2 is the heavy contamination case with outlier pattern (a). For the RMSPE, our method outperformed other comparative methods except in the case (p, ϵ, ρ) = (100, 0.3, 0.2) for sLTS with $\alpha = 0.5$. The Lasso also presented a worse performance and furthermore the sLTS with $\alpha = 0.75$ showed the worst performance due to a lack of truncation. For the TPR and TNR, our method showed the best performance.

Table 1: The outlier pattern (a) with $p = 100,\ 200$, $\epsilon = 0.1$ and $\rho = 0.2,\ 0.5$

	$p = 100$, $\epsilon = 0.1$, $\rho = 0.2$				$p = 100$, $\epsilon = 0.1$, $\rho = 0.5$			
Methods	RMSPE	MSE	TPR	TNR	RMSPE	MSE	TPR	TNR
Lasso	3.04	9.72×10^{-2}	0.936	0.909	3.1	1.05×10^{-1}	0.952	0.918
RLARS	0.806	6.46×10^{-3}	0.936	0.949	0.718	6.7×10^{-3}	0.944	0.962
sLTS($\alpha = 0.85$,80grids)	0.626	1.34×10^{-3}	1.0	0.964	0.599	1.05×10^{-3}	1.0	0.966
sLTS($\alpha = 0.75$,80grids)	0.651	1.71×10^{-3}	1.0	0.961	0.623	1.33×10^{-3}	1.0	0.961
sLTS($\alpha = 0.65$,80grids)	0.685	2.31×10^{-3}	1.0	0.957	0.668	1.76×10^{-3}	1.0	0.961
sparse γ-linear reg(γ=0.1)	0.557	6.71×10^{-4}	1.0	0.966	0.561	6.99×10^{-4}	1.0	0.965
sparse γ-linear reg(γ=0.5)	0.575	8.25×10^{-4}	1.0	0.961	0.573	9.05×10^{-4}	1.0	0.959
	$p = 200$, $\epsilon = 0.1$, $\rho = 0.2$				$p = 200$, $\epsilon = 0.1$, $\rho = 0.5$			
Methods	RMSPE	MSE	TPR	TNR	RMSPE	MSE	TPR	TNR
Lasso	3.55	6.28×10^{-2}	0.904	0.956	3.37	6.08×10^{-2}	0.928	0.961
RLARS	0.88	3.8×10^{-3}	0.904	0.977	0.843	4.46×10^{-3}	0.9	0.986
sLTS($\alpha = 0.85$,80grids)	0.631	7.48×10^{-4}	1.0	0.972	0.614	5.77×10^{-4}	1.0	0.976
sLTS($\alpha = 0.75$,80grids)	0.677	1.03×10^{-3}	1.0	0.966	0.632	7.08×10^{-4}	1.0	0.973
sLTS($\alpha = 0.65$,80grids)	0.823	2.34×10^{-3}	0.998	0.96	0.7	1.25×10^{-3}	1.0	0.967
sparse γ-linear reg($\gamma = 0.1$)	0.58	4.19×10^{-4}	1.0	0.981	0.557	3.71×10^{-4}	1.0	0.977
sparse γ-linear reg($\gamma = 0.5$)	0.589	5.15×10^{-4}	1.0	0.979	0.586	5.13×10^{-4}	1.0	0.977

Table 2: The outlier pattern (a) with $p = 100,\ 200$, $\epsilon = 0.3$ and $\rho = 0.2,\ 0.5$

	$p = 100$, $\epsilon = 0.3$, $\rho = 0.2$				$p = 100$, $\epsilon = 0.3$, $\rho = 0.5$			
Methods	RMSPE	MSE	TPR	TNR	RMSPE	MSE	TPR	TNR
Lasso	8.07	6.72×10^{-1}	0.806	0.903	8.1	3.32×10^{-1}	0.8	0.952
RLARS	2.65	1.54×10^{-1}	0.75	0.963	2.09	1.17×10^{-1}	0.812	0.966
sLTS($\alpha = 0.75$,80grids)	10.4	2.08	0.886	0.709	11.7	2.36	0.854	0.67
sLTS($\alpha = 0.65$,80grids)	2.12	3.66×10^{-1}	0.972	0.899	2.89	5.13×10^{-1}	0.966	0.887
sLTS($\alpha = 0.5$,80grids)	1.37	1.46×10^{-1}	0.984	0.896	1.53	1.97×10^{-1}	0.976	0.909
sparse γ-linear reg(γ=0.1)	1.13	9.16×10^{-2}	0.964	0.97	0.961	5.38×10^{-2}	0.982	0.977
sparse γ-linear reg(γ=0.5)	1.28	1.5×10^{-1}	0.986	0.952	1.00	8.48×10^{-2}	0.988	0.958
	$p = 200$, $\epsilon = 0.3$, $\rho = 0.2$				$p = 200$, $\epsilon = 0.3$, $\rho = 0.5$			
Methods	RMSPE	MSE	TPR	TNR	RMSPE	MSE	TPR	TNR
Lasso	8.11	3.4×10^{-1}	0.77	0.951	8.02	6.51×10^{-1}	0.81	0.91
RLARS	3.6	1.7×10^{-1}	0.71	0.978	2.67	1.02×10^{-1}	0.76	0.984
sLTS($\alpha = 0.75$,80grids)	11.5	1.16	0.738	0.809	11.9	1.17	0.78	0.811
sLTS($\alpha = 0.65$,80grids)	3.34	3.01×10^{-1}	0.94	0.929	4.22	4.08×10^{-1}	0.928	0.924
sLTS($\alpha = 0.5$,80grids)	4.02	3.33×10^{-1}	0.892	0.903	4.94	4.44×10^{-1}	0.842	0.909
sparse γ-linear reg($\gamma = 0.1$)	2.03	1.45×10^{-1}	0.964	0.924	3.2	2.86×10^{-1}	0.94	0.936
sparse γ-linear reg($\gamma = 0.5$)	1.23	7.69×10^{-2}	0.988	0.942	3.13	2.98×10^{-1}	0.944	0.94

Table 3 is the low contamination case with outlier pattern (b). For the RMSPE, our method outperformed other comparative methods (the oracle value of the RMSPE is 0.5). For the TPR and TNR, the sLTS showed a similar performance to our method. The Lasso presented the worst performance, because it is sensitive to outliers. Table 4 is the heavy contamination case with outlier pattern (b). For the RMSPE, our method outperformed

other comparative methods. The sLTS with $\alpha = 0.5$ showed the worst performance. For the TPR and TNR, it seems that our method showed the best performance.

Table 3: The outlier pattern (b) with $p = 100,\ 200$, $\epsilon = 0.1$ and $\rho = 0.2,\ 0.5$

	$p = 100$, $\epsilon = 0.1$, $\rho = 0.2$				$p = 100$, $\epsilon = 0.1$, $\rho = 0.5$			
Methods	RMSPE	MSE	TPR	TNR	RMSPE	MSE	TPR	TNR
Lasso	2.48	5.31×10^{-2}	0.982	0.518	2.84	5.91×10^{-2}	0.98	0.565
RLARS	0.85	6.58×10^{-3}	0.93	0.827	0.829	7.97×10^{-3}	0.91	0.885
sLTS($\alpha = 0.85$,80grids)	0.734	5.21×10^{-3}	0.998	0.964	0.684	3.76×10^{-3}	1.0	0.961
sLTS($\alpha = 0.75$,80grids)	0.66	1.78×10^{-3}	1.0	0.975	0.648	1.59×10^{-3}	1.0	0.961
sLTS($\alpha = 0.65$,80grids)	0.734	2.9×10^{-3}	1.0	0.96	0.66	1.74×10^{-3}	1.0	0.962
sparse γ-linear reg(γ=0.1)	0.577	8.54×10^{-4}	1.0	0.894	0.545	5.44×10^{-4}	1.0	0.975
sparse γ-linear reg(γ=0.5)	0.581	7.96×10^{-4}	1.0	0.971	0.546	5.95×10^{-4}	1.0	0.977
	$p = 200$, $\epsilon = 0.1$, $\rho = 0.2$				$p = 200$, $\epsilon = 0.1$, $\rho = 0.5$			
Methods	RMSPE	MSE	TPR	TNR	RMSPE	MSE	TPR	TNR
Lasso	2.39	2.57×10^{-2}	0.988	0.696	2.57	2.54×10^{-2}	0.944	0.706
RLARS	1.01	5.44×10^{-3}	0.896	0.923	0.877	4.82×10^{-3}	0.898	0.94
sLTS($\alpha = 0.85$,80grids)	0.708	1.91×10^{-3}	1.0	0.975	0.790	3.40×10^{-3}	0.994	0.97
sLTS($\alpha = 0.75$,80grids)	0.683	1.06×10^{-4}	1.0	0.975	0.635	7.40×10^{-4}	1.0	0.977
sLTS($\alpha = 0.65$,80grids)	1.11	1.13×10^{-2}	0.984	0.956	0.768	2.60×10^{-3}	0.998	0.968
sparse γ-linear reg($\gamma = 0.1$)	0.603	5.71×10^{-4}	1.0	0.924	0.563	3.78×10^{-3}	1.0	0.979
sparse γ-linear reg($\gamma = 0.5$)	0.592	5.04×10^{-4}	1.0	0.982	0.566	4.05×10^{-3}	1.0	0.981

Table 4: The outlier pattern (b) with $p = 100,\ 200$, $\epsilon = 0.3$ and $\rho = 0.2,\ 0.5$

	$p = 100$, $\epsilon = 0.3$, $\rho = 0.2$				$p = 100$, $\epsilon = 0.3$, $\rho = 0.5$			
Methods	RMSPE	MSE	TPR	TNR	RMSPE	MSE	TPR	TNR
Lasso	2.81	6.88×10^{-2}	0.956	0.567	3.13	7.11×10^{-2}	0.97	0.584
RLARS	2.70	7.69×10^{-2}	0.872	0.789	2.22	6.1×10^{-2}	0.852	0.855
sLTS($\alpha = 0.75$,80grids)	3.99	1.57×10^{-1}	0.856	0.757	4.18	1.54×10^{-1}	0.878	0.771
sLTS($\alpha = 0.65$,80grids)	3.2	1.46×10^{-1}	0.888	0.854	2.69	1.08×10^{-1}	0.922	0.867
sLTS($\alpha = 0.5$,80grids)	6.51	4.62×10^{-1}	0.77	0.772	7.14	5.11×10^{-1}	0.844	0.778
sparse γ-linear reg(γ=0.1)	1.75	3.89×10^{-2}	0.974	0.725	1.47	2.66×10^{-2}	0.976	0.865
sparse γ-linear reg(γ=0.5)	1.68	3.44×10^{-2}	0.98	0.782	1.65	3.58×10^{-2}	0.974	0.863
	$p = 200$, $\epsilon = 0.3$, $\rho = 0.2$				$p = 200$, $\epsilon = 0.3$, $\rho = 0.5$			
Methods	RMSPE	MSE	TPR	TNR	RMSPE	MSE	TPR	TNR
Lasso	2.71	3.32×10^{-2}	0.964	0.734	2.86	3.05×10^{-2}	0.974	0.728
RLARS	3.03	4.59×10^{-2}	0.844	0.876	2.85	4.33×10^{-2}	0.862	0.896
sLTS($\alpha = 0.75$,80grids)	3.73	7.95×10^{-2}	0.864	0.872	4.20	8.17×10^{-2}	0.878	0.87
sLTS($\alpha = 0.65$,80grids)	4.45	1.23×10^{-1}	0.85	0.886	3.61	8.95×10^{-2}	0.904	0.908
sLTS($\alpha = 0.5$,80grids)	9.05	4.24×10^{-1}	0.66	0.853	8.63	3.73×10^{-1}	0.748	0.864
sparse γ-linear reg($\gamma = 0.1$)	1.78	1.62×10^{-2}	0.994	0.731	1.82	1.62×10^{-2}	0.988	0.844
sparse γ-linear reg($\gamma = 0.5$)	1.79	1.69×10^{-2}	0.988	0.79	1.77	1.51×10^{-2}	0.996	0.77

5 Real data analyses

In this section, we use two real datasets to compare our method with comparative methods in real data analysis. We show the best result of comparative methods among some parameter situations (e.g., Sec 4.3), unless otherwise noted.

5.1 NCI-60 cancer cell panel

We applied our method and comparative methods to regress protein expression on gene expression data at the cancer cell panel of the National Cancer Institute. Experimental conditions were set in the same way as in Alfons et al. (2013) as follows. The gene expression data were obtained with an Affymetrix HG-U133A chip and normalized GCRMA method, resulting in a set of $p = 22,283$ explanatory variables. The protein expressions based on 162 antibodies were acquired via reverse-phase protein lysate arrays and $\log_2$ transformed. One observation had to be removed since all values were missing in the gene expression data, reducing the number of observations to $n = 59$. Then, the KRT18 antibody was selected as the response variable because it had the largest MAD among 162 antibodies, i.e., KRT18 may include a large number of outliers. Both the protein expressions and the gene expression data can be downloaded via the web application CellMiner (`http://discover.nci.nih.gov/cellminer/`). As a measure of prediction performance, the root trimmed mean squared prediction error (RTMSPE) was computed via leave-one-out cross-validation given by

$$\text{RTMSPE} = \sqrt{\frac{1}{h}\sum_{i=1}^{h}(e)^2_{[i:n]}}, \quad e^2 = \left(\left(y_1 - x_1^T\hat{\beta}^{[-1]}\right)^2, \ldots, \left(y_n - x_n^T\hat{\beta}^{[-n]}\right)^2\right),$$

where $(e)^2_{[1:n]} \leq \cdots \leq (e)^2_{[n:n]}$ are the order statistics of e^2 and $h = \lfloor (n+1)0.75 \rfloor$. The choice of h is important because it is preferable for estimating prediction performance that trimmed squares does not include outliers. We set h in the same way as in Alfons et al. (2013), because the sLTS detected 13 outliers in Alfons et al. (2013). In this experiment, we used the estimate of the RANSAC algorithm as an initial point instead of the sLTS because the sLTS required high computational cost in such a high dimensional data.

Table 5: Root trimmed mean squared prediction error (RTMSPE) for protein expressions based on the KRT18 antibody (NCI-60 cancer cell panel data), computed from leave-one-out cross validation

Methods	RTMSPE	# selected variables
Lasso	1.058	52
RLARS	0.936	18
sLTS	0.721	33
sparse γ-linear reg($\gamma = 0.1$)	0.679	29
sparse γ-linear reg($\gamma = 0.5$)	0.700	30

Table 5 shows that our method outperformed other comparative methods for the RTMSPE in high dimensional data. Our method presented the smallest RTMSPE with the second smallest number of explanatory variables. The RLARS presented the smallest number of explanatory variables, but a much larger RTMSPE than our method.

5.2 Protein homology dataset

We applied our method and comparative methods to protein sequences dataset used for KDD-Cup 2004. Experimental conditions were set in the same way as in Khan et al. (2007) as follows. The whole dataset consists of $n = 145,751$ protein sequences which has 153 blocks corresponding to native protein. Each data point in particular block is a candidate homologous protein. There were 75 variables in the data set: the block number (categorical) and 74 measurements of protein features. The first protein feature was used as the response variable. Then, 5 blocks with a total of $n = 4,141$ protein sequences were selected because they contained the highest proportions of homologous proteins (and hence the highest proportions of potential outliers). The data of each block was split into two almost equal parts to get a training sample of size $n_{tra} = 2,072$ and a test sample of size $n_{test} = 2,069$. The number of explanatory variables was $p = 77$, consisting of 4 block indicators (variables 1-4) and 73 features. The whole protein, train and test dataset can be downloaded from `http://users.ugent.be/~svaelst/software/RLARS.html`. As a measure of prediction performance, the root trimmed mean squared prediction error (RTMSPE) was computed for the test sample given by

$$\text{RTMSPE} = \sqrt{\frac{1}{h}\sum_{i=1}^{h}(e)^2_{[i:n_{test}]}}, \quad e^2 = \left(\left(y_1 - {x_1}^T\hat{\beta}\right)^2, \ldots, \left(y_{n_{test}} - x^T_{n_{test}}\hat{\beta}\right)^2\right),$$

where $(e)^2_{[1:n_{test}]} \leq \cdots \leq (e)^2_{[n_{test}:n_{test}]}$ are the order statistics of e^2 and $h = \lfloor (n_{test}+1)0.99 \rfloor$, $\lfloor (n_{test}+1)0.95 \rfloor$ or $\lfloor (n_{test}+1)0.9 \rfloor$. In this experiment, we used the estimate of the sLTS as an initial point.

Table 6: Root trimmed mean squared prediction error in the protein test set

Methods	trimming fraction 1%	5%	10%	# selected variables
Lasso	10.697	9.66	8.729	22
RLARS	10.473	9.435	8.527	27
sLTS	10.614	9.52	8.575	21
sparse γ-linear reg($\gamma = 0.1$)	10.461	9.403	8.481	44
sparse γ-linear reg($\gamma = 0.5$)	10.463	9.369	8.419	42

Table 6 shows that our method outperformed other comparative methods for the RTMSPE. Our method presented the smallest RTMSPE with the largest number of explanatory variables. It might seem that other methods gave the smaller number of explanatory variables than necessary.

Appendix

Proof of Theorem 2.1. For two non-negative functions $r(x, y)$ and $u(x, y)$ and probability density function $g(x)$, it follows from Hölder's inequality that

$$\int r(x,y)u(x,y)g(x)dxdy \leq \left(\int r(x,y)^\alpha g(x)dxdy \right)^{\frac{1}{\alpha}} \left(\int u(x,y)^\beta g(x)dxdy \right)^{\frac{1}{\beta}},$$

where α and β are positive constants and $\frac{1}{\alpha}+\frac{1}{\beta} = 1$. The equality holds if and only if $r(x, y)^\alpha = \tau u(x, y)^\beta$ for a positive constant τ. Let $r(x, y) = g(y|x)$, $u(x, y) = f(y|x)^\gamma$, $\alpha = 1 + \gamma$ and $\beta = \frac{1+\gamma}{\gamma}$. Then it holds that

$$\begin{aligned}
&\int \left(\int g(y|x) f(y|x)^\gamma dy \right) dg(x) \\
&\leq \left\{ \int \left(\int g(y|x)^{1+\gamma} dy \right) dg(x) \right\}^{\frac{1}{1+\gamma}} \left\{ \int \left(\int f(y|x)^{1+\gamma} dy \right) dg(x) \right\}^{\frac{\gamma}{1+\gamma}}.
\end{aligned}$$

The equality holds if and only if $g(y|x)^{1+\gamma} = \tau (f(y|x)^\gamma)^{\frac{1+\gamma}{\gamma}}$, i.e. $g(y|x) = f(y|x)$ because $g(y|x)$ and $f(y|x)$ are conditional probability density functions. The properties (i), (ii) follow from this inequality, the equality condition and the definition of $D_\gamma(g(y|x), f(y|x); g(x))$.

Let us prove the property (iii). Suppose that γ is sufficiently small. Then, it holds that $f^\gamma = 1 + \gamma \log f + O(\gamma^2)$. The γ-divergence for regression is expressed by

$$
\begin{aligned}
&D_\gamma(g(y|x), f(y|x); g(x)) \\
&= \frac{1}{\gamma(1+\gamma)} \log \int \left\{ \int g(y|x)(1 + \gamma \log g(y|x) + O(\gamma^2))dy \right\} g(x)dx \\
&\quad - \frac{1}{\gamma} \log \int \left\{ \int g(y|x)(1 + \gamma \log f(y|x) + O(\gamma^2))dy \right\} g(x)dx \\
&\quad \frac{1}{1+\gamma} \log \int \left\{ \int f(y|x)(1 + \gamma \log f(y|x) + O(\gamma^2))dy \right\} g(x)dx \\
&= \frac{1}{\gamma(1+\gamma)} \log \left\{ 1 + \gamma \int \left(\int g(y|x) \log g(y|x) dy \right) g(x)dx + O(\gamma^2) \right\} \\
&\quad - \frac{1}{\gamma} \log \left\{ 1 + \gamma \int \left(\int g(y|x) \log f(y|x) dy \right) g(x)dx + O(\gamma^2) \right\} \\
&\quad \frac{1}{1+\gamma} \log \left\{ 1 + \gamma \int \left(\int f(y|x) \log f(y|x) dy \right) g(x)dx + O(\gamma^2) \right\} \\
&= \frac{1}{(1+\gamma)} \int \left(\int g(y|x) \log g(y|x) dy \right) g(x)dx \\
&\quad - \int \left(\int g(y|x) \log f(y|x) dy \right) g(x)dx + O(\gamma) \\
&= \int D_{KL}(g(y|x), f(y|x)) g(x)dx + O(\gamma).
\end{aligned}
$$

□

Proof of Theorem 2.2. We see that

$$
\begin{aligned}
&\int \left(\int g(y|x) f(y|x;\theta)^\gamma dy \right) g(x)dx \\
&= \int \left(\int \{(1-\epsilon) f(y|x;\theta^*) + \epsilon \delta(y|x)\} f(y|x;\theta)^\gamma dy \right) g(x)dx \\
&= (1-\epsilon) \left\{ \int \left(\int f(y|x;\theta^*) f(y|x;\theta)^\gamma dy \right) g(x)dx \right\} + \epsilon \left\{ \int \left(\int \delta(y|x) f(y|x;\theta)^\gamma dy \right) g(x)dx \right\}.
\end{aligned}
$$

It follows from the assumption $\epsilon < \frac{1}{2}$ that

$$
\begin{aligned}
\left\{ \epsilon \int \left(\int \delta(y|x) f(y|x;\theta)^\gamma dy \right) g(x)dx \right\}^{\frac{1}{\gamma}} &< \left\{ \frac{1}{2} \int \left(\int \delta(y|x) f(y|x;\theta)^\gamma dy \right) g(x)dx \right\}^{\frac{1}{\gamma}} \\
&< \left\{ \int \left(\int \delta(y|x) f(y|x;\theta)^\gamma dy \right) g(x)dx \right\}^{\frac{1}{\gamma}} = \nu_{f_\theta,\gamma}.
\end{aligned}
$$

Hence,

$$\int \left(\int g(y|x) f(y|x;\theta)^\gamma dy \right) g(x) dx$$
$$= (1-\epsilon) \left\{ \int \left(\int f(y|x;\theta^*) f(y|x;\theta)^\gamma dy \right) g(x) dx \right\} + O\left(\nu_{f_\theta,\gamma}^\gamma\right).$$

Therefore, it holds that

$$\begin{aligned}
&d_\gamma(g(y|x), f(y|x;\theta); g(x)) \\
&= -\frac{1}{\gamma} \log \int \left(\int g(y|x) f(y|x;\theta)^\gamma dy \right) g(x) dx + \frac{1}{1+\gamma} \log \int \left(\int f(y|x;\theta)^{1+\gamma} dy \right) g(x) dx \\
&= -\frac{1}{\gamma} \log(1-\epsilon) - \frac{1}{\gamma} \log \int \left(\int f(y|x;\theta^*) f(y|x;\theta)^\gamma dy \right) g(x) dx + O\left(\nu_{f_\theta,\gamma}^\gamma\right) \\
&\qquad\qquad + \frac{1}{1+\gamma} \log \int \left(\int f(y|x;\theta)^{1+\gamma} dy \right) g(x) dx \\
&= -\frac{1}{\gamma} \log(1-\epsilon) + d_\gamma(f(y|x;\theta^*), f(y|x;\theta); g(x)) + O\left(\nu_{f_\theta,\gamma}^\gamma\right).
\end{aligned}$$

Then, it follows that

$$\begin{aligned}
&D_\gamma(g(y|x), f(y|x;\theta); g(x)) - D_\gamma(g(y|x), f(y|x;\theta^*); g(x)) - D_\gamma(f(y|x;\theta^*), f(y|x;\theta); g(x)) \\
&= \{-d_\gamma(g(y|x), g(y|x); g(x)) + d_\gamma(g(y|x), f(y|x;\theta); g(x))\} \\
&\quad - \{-d_\gamma(g(y|x), g(y|x); g(x)) + d_\gamma(g(y|x), f(y|x;\theta^*); g(x))\} \\
&\quad\quad - \{-d_\gamma(f(y|x;\theta^*), f(y|x;\theta^*); g(x)) + d_\gamma(f(y|x;\theta^*), f(y|x;\theta); g(x))\} \\
&= d_\gamma(g(y|x), f(y|x;\theta); g(x)) - d_\gamma(f(y|x;\theta^*), f(y|x;\theta); g(x)) \\
&\quad\quad - d_\gamma(g(y|x), f(y|x;\theta^*); g(x)) + d_\gamma(f(y|x;\theta^*), f(y|x;\theta^*); g(x)) \\
&= O\left(\nu^\gamma\right).
\end{aligned}$$

□

Proof of Theorem 2.3. We see that

$$\begin{aligned}
&\int \left(\int g(y|x) f(y|x;\theta)^\gamma dy \right) g(x) dx \\
&= \int \left(\int \{(1-\epsilon(x)) f(y|x;\theta^*) + \epsilon(x)\delta(y|x)\} f(y|x;\theta)^\gamma dy \right) g(x) dx \\
&= \left\{ \int \left(\int f(y|x;\theta^*) f(y|x;\theta)^\gamma dy \right) (1-\epsilon(x)) g(x) dx + \int \left(\int \delta(y|x) f(y|x;\theta)^\gamma dy \right) \epsilon(x) g(x) dx \right\}.
\end{aligned}$$

It follows from the assumption $\epsilon(x) < \frac{1}{2}$ that

$$\begin{aligned}\left\{\int\left(\int \delta(y|x)f(y|x;\theta)^\gamma dy\right)\epsilon(x)g(x)dx\right\}^{\frac{1}{\gamma}} &< \left\{\int\left(\int \delta(y|x)f(y|x;\theta)^\gamma dy\right)\frac{g(x)}{2}dx\right\}^{\frac{1}{\gamma}} \\ &< \left\{\int\left(\int \delta(y|x)f(y|x;\theta)^\gamma dy\right)g(x)dx\right\}^{\frac{1}{\gamma}} \\ &= \nu_{f_{\theta,\gamma}}.\end{aligned}$$

Hence,

$$\begin{aligned}&\int\left(\int g(y|x)f(y|x;\theta)^\gamma dy\right)g(x)dx \\ &\quad= \left\{\int\left(\int f(y|x;\theta^*)f(y|x;\theta)^\gamma dy\right)(1-\epsilon(x))g(x)dx\right\} + O(\nu^\gamma_{f_{\theta,\gamma}}).\end{aligned}$$

Therefore, it holds that

$$\begin{aligned}&d_\gamma(g(y|x), f(y|x;\theta); g(x)) \\ &= -\frac{1}{\gamma}\log\int\left(\int g(y|x)f(y|x;\theta)^\gamma dy\right)g(x)dx + \frac{1}{1+\gamma}\log\int\left(\int f(y|x;\theta)^{1+\gamma}dy\right)g(x)dx \\ &= -\frac{1}{\gamma}\log\left\{\int\left(\int f(y|x;\theta^*)f(y|x;\theta)^\gamma dy\right)(1-\epsilon(x))g(x)dx\right\} + O(\nu^\gamma_{f_{\theta,\gamma}}) \\ &\quad + \frac{1}{1+\gamma}\log\int\left(\int f(y|x;\theta)^{1+\gamma}dy\right)g(x)dx \\ &= d_\gamma(f(y|x;\theta^*), f(y|x;\theta); (1-\epsilon(x))g(x)) + O(\nu^\gamma_{f_{\theta,\gamma}}) \\ &\quad - \frac{1}{1+\gamma}\log\int\left(\int f(y|x;\theta)^{1+\gamma}dy\right)(1-\epsilon(x))g(x)dx \\ &\qquad + \frac{1}{1+\gamma}\log\int\left(\int f(y|x;\theta)^{1+\gamma}dy\right)g(x)dx \\ &= d_\gamma(f(y|x;\theta^*), f(y|x;\theta); (1-\epsilon(x))g(x)) + O(\nu^\gamma_{f_{\theta,\gamma}}) - \frac{1}{1+\gamma}\log\left\{1-\int\epsilon(x)g(x)dx\right\}.\end{aligned}$$

Then, it follows that

$$\begin{aligned}&D_\gamma(g(y|x), f(y|x;\theta); g(x)) \\ &- D_\gamma(g(y|x), f(y|x;\theta^*); g(x)) - D_\gamma(f(y|x;\theta^*), f(y|x;\theta); (1-\epsilon(x))g(x)) \\ &= \{-d_\gamma(g(y|x), g(y|x); g(x)) + d_\gamma(g(y|x), f(y|x;\theta); g(x))\} \\ &\quad - \{-d_\gamma(g(y|x), g(y|x); g(x)) + d_\gamma(g(y|x), f(y|x;\theta^*); g(x))\}\end{aligned}$$

$$
\begin{aligned}
&\quad - \{-d_\gamma(f(y|x;\theta^*), f(y|x;\theta^*); (1-\epsilon(x))g(x)) + d_\gamma(f(y|x;\theta^*), f(y|x;\theta); (1-\epsilon(x))g(x))\} \\
&= d_\gamma(g(y|x), f(y|x;\theta); g(x)) - d_\gamma(f(y|x;\theta^*), f(y|x;\theta); (1-\epsilon(x))g(x)) \\
&\quad - d_\gamma(g(y|x), f(y|x;\theta^*); g(x)) + d_\gamma(f(y|x;\theta^*), f(y|x;\theta^*); (1-\epsilon(x))g(x)) \\
&= O\left(\nu^\gamma\right).
\end{aligned}
$$

□

References

A. Alfons, C. Croux, and S. Gelper. Sparse least trimmed squares regression for analyzing high-dimensional large data sets. *Ann. Appl. Stat.*, 7(1): 226–248, 2013.

A. Basu, I. R. Harris, N. L. Hjort, and M. C. Jones. Robust and efficient estimation by minimising a density power divergence. *Biometrika*, 85(3): 549–559, 1998.

A. Basu, H. Shioya, and C. Park. *Statistical Inferefnce:The Minimum Distance Approach*. Boca Raton:Taylor & Francis, 2010.

B. Efron, T. Hastie, I. Johnstone, and R. Tibshirani. Least angle regression. *Ann. Statist.*, 32(2):407–499, 2004.

J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani. Pathwise coordinate optimization. *Ann. Appl. Stat.*, 1(2):302–332, 2007.

H. Fujisawa and S. Eguchi. Robust parameter estimation with a small bias against heavy contamination. *J. Multivar. Anal.*, 99(9):2053–2081, 2008.

T. Hastie, R. Tibshirani, and J. Friedman. *The Elements of Statistical Learning: Data Mining, Inference and Prediction*. Springer, 2010.

D. R. Hunter and K. Lange. A tutorial on MM algorithms. *Am. Stat.*, 58 (1):30–37, 2004.

M. C. Jones, N. L. Hjort, I. R. Harris, and A. Basu. A comparison of related density-based minimum divergence estimators. *Biometrika*, 88(3):865–873, 2001.

T. Kanamori and H. Fujisawa. Robust estimation under heavy contamination using unnormalized models. *Biometrika*, 102(3):559–572, 2015.

J. A. Khan, S. Van Aelst, and R. H. Zamar. Robust linear model selection based on least angle regression. *J. Am. Statist. Assoc.*, 102:1289–1299, 2007.

R. A. Maronna, D. R. Martin, and V. J. Yohai. *Robust Statistics: Theory and Methods*. John Wiley and Sons, 2006.

P. J. Rousseeuw. Least median of squares regression. *J. Am. Statist. Assoc.*, 79(388):871–880, 1984.

R. Tibshirani. Regression shrinkage and selection via the lasso. *J. R. Statist. Soc. B*, 58:267–288, 1996.

M. P. Windham. Robustifying model fitting. *J. R. Statist. Soc. B*, 57(3): 599–609, 1995.

H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *J. R. Statist. Soc. B*, 67:301–320, 2005.